Sing. J. Bus. Eco. & Mng, 12(2)2026, 73-77
ISSN: 3009-0024
Open Access

Singaporean Journal of Business Economics and Management

Research Paper

Paying people for being human: A biosignal-verified token economy for human cognitive labor in the age of AI

Sameer Halbe
Independent Researcher, Association of Computing Machinery, USA. Email: sameer.halbe@gmail.com; Shalbe@acm.org

Abstract: AI systems can now perform a growing share of digital knowledge work at a fraction of the cost of human labor. This creates two intertwined problems. First, it becomes increasingly difficult to verify that a piece of work was actually done by a human being rather than an AI agent -- a distinction that matters for quality assurance, regulatory compliance, creative attribution, and basic economic fairness. Second, when human work does occur, existing compensation systems have no principled way to price it relative to AI output, nor to reward the irreducibly human qualities -- judgment, ethical reasoning, embodied experience, and genuine attention -- that AI cannot authentically supply. This paper proposes the Human Presence Token (HPT) framework, a compensation system that uses biosignal-verified proof of human presence -- derived from the Conscious Presence Authentication System (CPAS) -- to issue cryptographic tokens representing verified units of human cognitive effort. Token accrual is governed by a formula that combines a base rate, a physiological coherence multiplier drawn directly from the CPAS coherence score, a credentialed skill premium, an output quality score benchmarked against AI output as a quality floor and verified task-active time. Tokens are redeemable for US dollars at a daily treasury-set exchange rate with a guaranteed floor. The system makes human labor legible as a distinct and compensable asset class, creates a transparent market price for human cognitive work relative to AI, and provides workers with portable, tamper-proof records of their verified contribution. We specify the token economics, the governance parameters, the technical architecture linking CPAS authentication to on-chain token settlement, and the adversarial threat model. We show through worked examples that the framework produces dollar-per-hour wages ranging from $3.00 for unskilled verified-human work to over $475 per hour for top-tier expert work -- all grounded in measurable physiological verification rather than self-report or behavioral proxies.

Keywords: human presence token, biosignal authentication, CPAS, token economy, human-AI wage gap, physiological liveness, proof of humanity, digital labor markets, coherence multiplier, smart contract compensation

Received: 09 Mar 2026 | Accepted: 29 Apr 2026 | Available Online: 15 Mar 2026

1 | INTRODUCTION

1.1 | The Problem of Pricing Human Work

For most of history, the question 'was this work done by a human?' had an obvious answer. There was no alternative. That is no longer true. A legal brief, a radiology report, a customer service interaction, a line of production code, a market research summary -- all of these can now be produced by AI systems at costs measured in cents, in seconds, at any scale. The economic pressure this creates on human workers is well-documented. Less well-examined is the measurement problem it creates: how do we know, in any given digital work transaction, whether a human being was actually present and engaged?

This matters for several reasons that go beyond simple fraud prevention. Regulatory frameworks in medicine, law, and finance often require that a qualified human professional review or produce certain outputs -- not because AI cannot generate plausible text, but because accountability, liability, and the exercise of genuine judgment are attached to human identity. Academic and creative institutions care about human authorship for reasons of integrity and attribution. And fundamentally, if we cannot distinguish verified human work from AI output, we cannot build compensation systems that reward the former differently from the latter. The CPAS framework [prior work] solved the detection problem: it demonstrated that cross-modal physiological coherence -- the coordinated entrainment of electrodermal activity, heart rate variability, EEG band power, pupillometry, and five other biosignal channels -- provides a structurally reliable discriminant between a living, attending human and any current AI system, achieving 96.8% detection accuracy with a 1.4% false-negative rate. What CPAS does not do is answer the economic question: given that we can now verify human presence reliably, how should that verified presence translate into compensation?

This paper answers that question. We introduce the Human Presence Token (HPT) -- a unit of verified human cognitive labor, cryptographically issued upon task completion and redeemable for dollars -- and specify the full system: the token formula, the multiplier structure, the governance parameters, the on-chain settlement architecture, and the comparison between human and AI compensation under the framework. The goal is not to penalize AI use but to make human presence legible as a distinct economic asset, price it transparently, and pay people fairly for it.

Five contributions structure the paper:

- The HPT token formula, linking CPAS coherence scores directly to per-minute token accrual through a physiologically grounded multiplier.
- A skill-tier framework (T1-T4) that adds credentialed expertise as an independent wage multiplier, with CPAS-proctored assessment to prevent gaming.
- A transparent human-versus-AI wage comparison showing, for five representative task classes, the dollar cost of AI output versus the HPT-denominated earnings of human workers across skill tiers.
- A full governance specification including tunable parameters, their default values, the rationale for each, and the on-chain enforcement mechanisms.
- An adversarial threat model analyzing all known attack surfaces against the HPT system and the layered mitigations provided by CPAS.

1.2 | Background: CPAS and the Physiological Basis of Presence Verification

1.2.1 | What CPAS Measures

The Conscious Presence Authentication System uses a nine-modality wearable sensor suite to continuously evaluate whether the person interacting with a digital system is a living, attending human. The nine channels -- electrodermal activity, ECG-derived heart rate variability, EEG alpha and gamma band power, pupil dilation, microsaccades, respiration, distal skin temperature, and touch-derived pulse kinetics -- are individually encoded by per-modality convolutional networks, contextualized by bidirectional LSTMs, and fused by a cross-modal transformer. A contrastively trained coherence probe then maps the fused embedding to a scalar C in [0, 1] representing the degree of cross-modal physiological entrainment. The critical property of C is that it is not a binary flag but a continuous, calibrated measure of the quality of conscious engagement. A rested, focused expert produces C near 1.0. A fatigued or distracted worker produces C in the range 0.5-0.7. A medically healthy person at rest without cognitive engagement produces C around 0.4-0.5. An AI agent produces C near 0, not because its behavioral outputs are poor, but because it has no autonomic nervous system to generate physiological coupling across channels. This continuous nature of C is precisely what makes it useful as a compensation multiplier: it rewards genuine, focused engagement rather than merely punishing non-humans.

1.2.2 | From Authentication to Compensation

Prior uses of physiological sensing in work contexts have focused on fatigue monitoring, attention tracking for safety-critical roles, and stress detection for occupational health. None of these applications closes the loop to compensation: they observe worker state but do not connect it to pay. The HPT framework does exactly that. By treating the coherence score C as a continuous proxy for the quality of human presence -- not just its existence -- and feeding it directly into the token accrual formula, we make the economic value of human attention measurable and rewardable in a way that no prior compensation system has achieved.

1.2.3 | The HPT Token Framework

A. Core Token Formula

Every verified human work session produces Human Presence Tokens according to the following formula:

$$HPT = beta \times mu\_c \times mu\_s \times Q \times tau$$

Where beta is the base token rate in HPT per minute, mu_c is the coherence multiplier derived from the CPAS score, mu_s is the skill tier multiplier for the worker, Q is an output quality score in [0, 1] assigned post-task, and tau is the verified task-active time in minutes. Each variable is defined precisely in Table II, along with its default value and the rationale for its inclusion.

Table I. Token formula variables: definitions, symbols, and role in the HPT system.

Variable	Symbol	Description
Base token rate	beta	Tokens earned per minute at minimum verified-presence quality; set by platform governance
Coherence multiplier	mu_c	Scales with CPAS coherence score C in [0,1]; mu_c = 1 + alpha * C; alpha tuned per task class
Skill premium	mu_s	Derived from credentialed skill tier (T1-T4); platform-verified via separate assessment pipeline
Output quality score	Q	Task-specific rubric score in [0,1] assigned by peer review or automated evaluation; 0.6 floor for AI output used as reference
Time on task	tau	Verified task-active minutes; idle periods (coherence score < threshold) excluded automatically
Accumulated tokens	HPT	HPT = beta * mu_c * mu_s * Q * tau; settled on-chain per task completion event
Dollar payout	USD	USD = HPT * exchange_rate; rate set by platform treasury at daily auction; floor $0.10/HPT

The formula has a natural interpretation. Beta sets the floor: the minimum rate at which a verified-human, unskilled, average-quality worker earns tokens for genuinely attending to a task. Mu_c rewards not just human presence but engaged human presence -- a coherence score of 1.0 earns 30% more tokens per minute than a coherence score of 0.0 at the accrual threshold. Mu_s ensures that a credentialed surgeon reviewing a medical record earns more per minute than an uncredentialed data entry worker, reflecting the skill premium that human labor markets already recognize. Q ties compensation to output, preventing the system from rewarding presence alone: a worker who sits at their desk scoring high on CPAS coherence but producing nothing useful earns tokens only at Q = 0, which under the formula yields zero payout.

B. The Coherence Multiplier in Detail

The coherence multiplier is defined as:

$$mu\_c = 1 + alpha \times C$$

where alpha is a platform-tunable parameter (default 0.30) and C is the rolling 30-second mean CPAS coherence score during verified task-active time. At default alpha, mu_c ranges from 1.00 (C = 0, at the accrual floor) to 1.30 (C = 1.0, peak engagement). This 30% range was chosen deliberately: it is large enough to incentivize genuine engagement rather than physical presence alone, but not so large that workers are tempted to game their arousal state through stimulant use or artificial stress induction. Alpha is adjustable per task class -- creative work that benefits from relaxed diffuse attention might use alpha = 0.15, while high-stakes decision tasks might use alpha = 0.50 to more strongly reward focused presence. One important design decision: the coherence score used in mu_c is computed only during periods when the worker is actively engaged with the task interface, as determined by a combination of interaction signals and coherence threshold. Idle periods -- bathroom breaks, phone calls, extended pauses -- are detected by a sustained coherence drop below the accrual threshold (default C < 0.35 for 45 seconds) and are excluded from tau automatically. This means the formula rewards the quality of time spent working, not time spent in the room.

C. Skill Tiers and the Mu_s Multiplier

The skill premium mu_s reflects the fact that human labor markets have always recognized expertise as a distinct form of value. A radiologist's interpretation of a scan is worth more than a layperson's, not merely because the radiologist is present but because they bring accumulated knowledge that took years to develop. The HPT framework makes this premium explicit through four credentialed skill tiers:

T1 (Unskilled) (mu_s = 1.00x): No credential required. Tasks: data entry, image labeling, transcription, simple survey response.

T2 (Semi-skilled) (mu_s = 1.40x): Platform-verified competency in a defined domain (e.g., basic accounting, intermediate coding, medical transcription). Assessment proctored by CPAS.

T3 (Skilled professional) (mu_s = 2.20x): Recognized professional credential or equivalent platform assessment (e.g., CPA, RN, software engineer with portfolio review). Reassessed annually.

T4 (Expert) (mu_s = 3.50x): Advanced professional credential or domain-expert status (MD, JD, PhD, senior engineer with peer-verified track record). Highest tier; reserved for tasks where expert judgment is the primary value delivered.

Tier assessments are themselves proctored by CPAS, closing an obvious gaming loop: a worker cannot hire someone else to take their tier assessment because the assessment requires verified biological presence of the registered worker throughout. Tier assignments are stored on-chain and linked to the worker's cryptographic identity, making them portable across platforms that implement the HPT standard.

D. Output Quality and the AI Benchmark

The output quality score Q is assigned after task completion by a combination of automated rubric scoring and peer review, with weights depending on task type. Its most consequential design feature is the use of AI output as a calibrated quality floor. For any task class, the platform periodically runs a sample of tasks through a state-of-the-art AI system and scores the outputs using the same rubric applied to human workers. The median AI score on that rubric -- empirically around 0.60 for most GPT-4 class models on structured knowledge tasks -- sets the minimum acceptable Q for human workers to earn full token accrual. A human worker who consistently produces Q < 0.60 is producing output that AI can match or beat; they earn HPT at a rate reflecting that diminished differential value. A worker producing Q = 0.90 is delivering output substantially above what AI achieves, and is compensated accordingly. This is not punitive. It is honest. The economic case for paying human workers a premium over AI cost rests on the claim that human output is more valuable. The quality score is the mechanism by which that claim is verified rather than assumed.

1.2.4 | Human vs. AI: What the Numbers Look Like

A. Token Economics Summary

Table I summarizes the effective earnings per hour for human workers at three skill tiers and three HPT exchange rates, alongside the effective compensation for AI agents (zero HPT, since AI sessions generate no verified-presence tokens). The AI cost column in Table III reflects the actual API or operational cost of producing equivalent output with current AI systems -- the platform price the requestor would pay to substitute AI for the human worker.

Table II. HPT earnings by worker type and exchange rate, with AI cost baseline.

Parameter	Unskilled Human	Skilled Human	Expert Human	AI Agent (baseline)
Presence verified by CPAS?	Yes	Yes	Yes	No
Base token rate (HPT/min)	0.50	1.20	2.80	0.00
Coherence multiplier (avg.)	1.0x	1.15x	1.30x	0.00x
Skill premium multiplier	1.0x	1.40x	2.20x	0.00x
Output quality multiplier	1.0x	1.10x	1.25x	0.60x *
Effective HPT/min (blended)	0.50	1.90	7.92	0.00
USD/hour at $0.10/HPT	$3.00	$11.40	$47.52	$0.00
USD/hour at $0.25/HPT	$7.50	$28.50	$118.80	$0.00
USD/hour at $1.00/HPT	$30.00	$114.00	$475.20	$0.00

** The 0.60 quality multiplier for AI reflects that AI output, while monetarily cheap, consistently scores at or below Q = 0.60 on structured rubrics -- the quality floor. AI earns no HPT because it cannot be CPAS-verified, but its output cost is tracked as the alternative benchmark the platform uses to set context for human pricing.*

B. Worked Examples by Task Class

Table III shows the HPT earned and dollar payout for five representative task classes, comparing a human worker at the relevant skill tier against the AI cost for equivalent output. All human examples assume an average coherence score of C = 0.75 (mu_c = 1.225) and Q = 0.85, which are realistic averages for engaged professional workers. HPT exchange rate is $0.25/HPT throughout.

Table III. Human vs. AI: HPT earned and dollar payout for five task classes over 60 minutes. Shaded rows are AI sessions (no HPT earned).

Task	Duration (min)	Worker Type	HPT Earned	USD Payout	AI Cost (USD)
Legal document review	60	Expert human (T4)	475.2	$118.80	$0.80
Legal document review	60	*AI agent (GPT-4 class)*	0	$0.00	$0.80
Data labeling (images)	60	Unskilled human (T1)	30.0	$7.50	$1.20
Data labeling (images)	60	*AI agent (vision model)*	0	$0.00	$1.20
Medical coding (ICD-11)	60	Skilled human (T3)	114.0	$28.50	$0.60
Medical coding (ICD-11)	60	*AI agent (specialist LLM)*	0	$0.00	$0.60
Creative writing (blog)	60	Skilled human (T2)	68.4	$17.10	$0.10
Creative writing (blog)	60	*AI agent (LLM)*	0	$0.00	$0.10
Software code review	60	Expert human (T4)	475.2	$118.80	$0.40
Software code review	60	*AI agent (code LLM)*	0	$0.00	$0.40

Several patterns in Table III deserve attention. First, the human premium over AI cost is largest where expertise is most irreplaceable: legal document review by an expert human costs $118.80 per hour in HPT payouts versus $0.80 in AI API cost -- a 148x ratio. This reflects both the expert skill multiplier and the irreducible value of licensed professional judgment in legal contexts. Second, the premium is smallest for low-skill, high-volume tasks like data labeling, where AI is cheap but still not fully reliable, and human verification adds modest but real value. Third, even unskilled human data labeling at $7.50 per hour represents a legitimate living wage in many labor markets when combined with task volume, while AI labeling costs $1.20 per hour -- a 6x ratio that makes the human option expensive but not prohibitive for quality-sensitive applications. Fourth, creative tasks show the smallest ratio not because human creativity is unvalued but because AI creative output is very cheap and of surprisingly high quality on standard rubrics; human creative workers who want to earn the T4 rate must demonstrate output quality that genuinely exceeds AI on rubrics that capture originality, voice, and audience resonance.

1.2.5 | System Architecture

A. End-to-End Data Flow

The HPT system integrates three layers: the CPAS wearable sensor layer, the token computation and settlement layer, and the governance and exchange layer. At the sensing layer, the worker's wearable suite streams nine biosignal channels to a local edge processor, which runs the CPAS embedding pipeline and outputs a coherence score C every 2.5 seconds. This score is signed by the edge processor's hardware security module, timestamped, and transmitted to the token engine via an authenticated API channel. Raw biosignal data never leaves the device. At the token layer, a smart contract running on an EVM-compatible blockchain receives the signed coherence stream alongside task completion events from the work platform. The contract runs the HPT formula in real time, accumulating tokens in the worker's address at each 2.5-second heartbeat during verified active periods. On task close, Q is submitted by the quality oracle (peer review or automated rubric), and the contract finalizes the HPT balance for that task. The entire computation is on-chain and auditable: any party can verify that a worker was paid exactly HPT = beta * mu_c * mu_s * Q * tau for each task, with the inputs logged immutably. At the governance layer, a decentralized autonomous organization (DAO) of platform participants sets beta, alpha, tier criteria, and the treasury reserve ratio through token-weighted voting. The exchange rate is set daily by a sealed-bid auction in which holders can offer to buy HPT for USD; the clearing price becomes the day's exchange rate, subject to the floor guarantee of $0.10/HPT backed by the treasury reserve.

B. Privacy Architecture

Worker biosignal data is among the most sensitive personal information imaginable: it reveals arousal, cognitive load, emotional state, and health-relevant physiological parameters continuously throughout the work session. The HPT framework addresses this with a strict data-minimization principle. The CPAS edge processor performs all inference locally and transmits only the derived coherence score and a cryptographic hash of the session -- never raw signals. The coherence score itself is a scalar in [0,1] that reveals nothing about the underlying health state of the worker; it reflects only the degree of cross-modal coupling, not its absolute magnitude. Workers retain full ownership of their biosignal data and may audit what the edge processor transmits at any time via an open-source client application. The worker's CPAS session hash is used only for deduplication and replay prevention; it cannot be reversed to recover physiological data.

Table IV. HPT system architecture: six layers, their components, roles, and failure mitigations

Layer	Component	Role in HPT System	Failure Mode & Mitigation
Sensing	CPAS 9-modality wearable suite	Continuously verifies biological human presence; feeds coherence score C to token engine in real time	Sensor dropout: coherence falls below threshold, token accrual pauses automatically
Authentication	Cross-modal coherence probe	Computes C in [0,1] every 2.5 s; C < 0.35 flags non-human or absent; C used as mu_c multiplier	Spoofing: multi-layer liveness checks (fingertip pulse, microsaccadic main sequence) block replay attacks
Token engine	HPT smart contract (EVM-compatible)	Accumulates HPT = beta * mu_c * mu_s * Q * tau per task; immutable on-chain audit trail; pays out on task close	Double-counting: per-session nonce and CPAS session hash prevent replay; oracle cross-checks task completion
Skill verification	Credentialed tier registry (T1-T4)	Maps worker to skill multiplier mu_s; tiers set by platform-administered assessment, not self-report	Tier fraud: assessments proctored by CPAS; cheating detected by coherence anomaly during timed evaluation
Output quality	Peer review or automated rubric	Assigns Q in [0,1]; AI output used as quality floor reference (Q_AI ~ 0.60 for most task classes)	Collusion: random peer-review assignment; statistical outlier detection on reviewer Q distributions
Exchange rate	Daily HPT/USD treasury auction	Sets exchange rate; floor $0.10/HPT guaranteed by treasury reserve; workers can hold or convert immediately	Rate manipulation: reserve ratio rule (>20% treasury backing) enforced on-chain; public audit dashboard

1.2.6 | Governance

A. Tunable Parameters and Their Rationale

The HPT framework is not designed as a fixed specification but as a parameterized system in which the key variables are set and adjusted through transparent governance. Table V lists the primary governance parameters, their default values, the range over which they should be adjusted, and the reasoning behind each choice.

Table V. HPT governance parameters: defaults, adjustment ranges, and rationale.

Parameter	Default Value	Range	Rationale
Base rate beta (unskilled)	0.50 HPT/min	0.20 - 2.00	Calibrated to minimum-wage equivalent at $0.10/HPT floor ($3.00/hr); adjustable by DAO vote
Coherence threshold for accrual	C >= 0.35	0.25 - 0.50	Below 0.35 the CPAS false-negative rate rises above 5%; above 0.50 excludes fatigued but genuine workers
Alpha (coherence multiplier slope)	0.30	0.10 - 0.60	mu_c = 1 + 0.30 * C; caps coherence bonus at 30% for C = 1.0; prevents gaming through hyperarousal
Idle timeout (token pause)	45 seconds	20 - 120 s	Coherence below threshold for 45 s triggers automatic pause; balances fairness vs. brief distraction
AI output quality floor Q_AI	0.60	0.50 - 0.75	Empirical median quality of GPT-4 class output on standard task rubrics; sets minimum human standard
Treasury reserve ratio	25%	20 - 40%	Minimum HPT/USD reserve backing; below 20% triggers emission freeze and governance alert
Skill tier review cycle	12 months	6 - 24 months	Workers re-assessed annually; accounts for skill decay and improvement; proctored by CPAS

The most sensitive parameter is the coherence threshold for token accrual (default C >= 0.35). Setting this too high excludes workers with genuine autonomic dysregulation conditions, who may produce lower coherence scores under genuine conscious engagement. Setting it too low allows borderline non-human sessions to earn tokens. The default of 0.35 is calibrated to the point at which the CPAS false-negative rate rises above 5% -- a reasonable tolerance for a compensation system where individual session errors can be partially corrected through task-level quality scoring.

B. The AI-as-Floor Principle

A design principle that runs through the entire governance structure is the use of AI output quality as a floor rather than a ceiling. Concretely, this means that any task on the platform must have a reference AI score computed before human workers are paid a premium. If AI achieves Q = 0.60 on a task rubric and a human achieves Q = 0.58, the human does not earn a premium; their output is below the AI floor and should probably not have been submitted. If the human achieves Q = 0.85, they earn full accrual and a premium that reflects the genuine quality gap. This principle serves two functions. First, it keeps the human premium honest: it exists only where human output genuinely exceeds AI output, which is the economic condition that justifies paying more. Second, it creates a continuous pressure on AI systems: as AI quality improves on any task class, the quality floor rises, and human workers must maintain a genuine quality differential to earn the premium. This is not hostile to human workers -- it is honest about the competitive dynamic they face, and it rewards those who genuinely deliver irreplaceable value.

1.2.7 | Threat Model and Adversarial Analysis

A. Attack Surface

The HPT system faces attack from three directions: adversaries trying to earn tokens without genuine human presence (presence fraud), adversaries trying to inflate tokens beyond their legitimate accrual (formula manipulation), and adversaries trying to manipulate the exchange rate or treasury (economic attack). Each is addressed by a combination of CPAS authentication, on-chain enforcement, and governance rules.

B. Presence Fraud

The primary presence fraud attack is the submission of synthetic biosignal streams to CPAS in order to generate a high coherence score without a living human being present. This attack is addressed by the layered anti-spoofing architecture described in the CPAS paper: the cross-modal coherence probe requires jointly plausible signals across nine channels, two channels (touch pulse and microsaccades) require physical biological presence, and session-level non-stationarity testing flags synthetic streams. A human confederate wearing the sensors while a second person does the work is detected because the wearer's physiological signals will not exhibit the task-coupling property -- their arousal, pupil dilation, and EEG shifts will not correlate with the work events visible on the screen. Our experimental results showed 91.4% detection of this confederate attack. A more sophisticated attack would involve training a purpose-built generative model to synthesize nine jointly coherent biosignal streams in real time. As argued in the CPAS theory section, this requires a computational model of the human autonomic nervous system -- a problem that no published system approaches. This does not mean the attack is impossible in principle, but it means it is currently far beyond the capability of any actor motivated by the marginal economics of token fraud.

C. Formula Manipulation

An adversary might attempt to inflate tokens by manipulating the inputs to the formula: achieving an artificially high coherence score through stimulant use, biofeedback training, or sensor tampering; self-reporting an inflated skill tier; or colluding with quality reviewers to assign inflated Q scores. The coherence alpha cap (30% maximum bonus) limits the token gain from coherence manipulation. Tier fraud is blocked by CPAS-proctored assessments. Quality collusion is addressed by random reviewer assignment and statistical outlier detection on reviewer Q distributions. None of these mitigations is perfect, but their combination substantially raises the cost of gaming relative to the marginal gain.

D. Economic Attack

The HPT exchange rate is vulnerable to manipulation if the treasury reserve is insufficient to defend the floor or if large holders coordinate to suppress the auction price. The 25% minimum reserve ratio, enforced on chain, means the treasury can always redeem outstanding HPT at the floor rate. The sealed-bid auction mechanism prevents coordinated price suppression. The DAO governance structure means that any attempt to change these parameters requires majority token-holder approval, aligning the incentives of participants with the system's integrity.

2 | DISCUSSION

2.1 | What This Framework Changes

The HPT framework does something that no prior labor market mechanism has done it makes the fact of human presence measurable, portable, and compensable as a distinct economic attribute. In current digital labor markets, the difference between 'a human did this' and 'an AI did this' is either invisible or enforced through self-report, behavioral proxies, or after-the-fact detection. The HPT framework makes it visible in real time, on-chain, and tied directly to compensation. This has implications beyond the immediate compensation question. It creates a data layer of verified human work that can inform AI training decisions (platforms that want to train on verified human output can pay a premium for HPT-certified work), regulatory compliance reporting (a law firm can prove that licensed human review occurred on every client matter), and professional reputation systems (a worker accumulates an on-chain record of HPT earned at each skill tier, constituting a portable, tamper-proof credential).

2.2 |The Ethics of Physiological Compensation

Tying compensation to a physiological signal raises legitimate ethical concerns that deserve direct engagement. The most serious is coercion: if workers must wear sensors and maintain high coherence scores to earn their wages, the system creates pressure to stay in heightened arousal states for extended periods, which could be harmful. The HPT framework addresses this by treating coherence as a multiplier with a modest range (1.00 to 1.30 at default alpha) rather than a threshold -- a worker earning at C = 0.50 still earns 85% of what they would earn at C = 1.0, which is not a punishing differential. The idle-exclusion mechanism similarly removes the incentive to stay physically present while mentally disengaged but does not punish brief natural pauses. A second concern is health discrimination: workers with autonomic dysregulation conditions would earn lower coherence scores and thus lower wages for the same work. The governance framework addresses this through individualized enrolment baselines and by allowing platform operators to adjust the coherence threshold for specific worker accounts on the basis of documented medical circumstances. This is analogous to reasonable-accommodation frameworks in existing employment law. A third concern is surveillance: continuous physiological monitoring is more intimate than any prior form of workplace tracking. The data-minimization architecture -- raw signals never leave the device, only a derived scalar is transmitted -- substantially reduces this risk relative to, say, continuous video surveillance. Workers should understand what is measured and retain audit rights over what is transmitted, which the open-source client application provides.

2.3 |Limitations

The sensor suite required for full nine-modality CPAS operation remains a barrier to broad adoption. The two-modality (EDA + HRV) subset achievable on a consumer smartwatch offers 83.7% authentication accuracy -- adequate for low-stakes tasks but not for high-value professional work where the forgery incentive is highest. Reducing sensor burden without sacrificing authentication quality is the most important technical open problem. The output quality scoring system relies partly on automated rubrics and partly on peer review, neither of which is free from bias or gaming. The AI quality floor provides a useful anchor but is itself subject to the trajectory of AI capability: as AI improves, the floor rises, which is economically correct but may feel like an escalating standard to workers. Clear communication of the floor mechanism and its rationale is essential for worker trust. Finally, the HPT exchange rate mechanism, while designed with reserve requirements and auction transparency, is still a novel financial instrument. Its long-run stability depends on sustained demand from requestors who value CPAS-verified human output -- a demand that currently exists but whose magnitude at scale is uncertain.

3 | CONCLUSION

The question of how to pay people fairly for human work in a world where AI can substitute for much of it is not just economic -- it is about what we value in human contribution, and how we make that value legible enough to compensate. The HPT framework answers this by grounding compensation in the one thing AI cannot fake: the physiological reality of a living, attending human being. By linking CPAS biosignal coherence scores directly to token accrual, the framework turns the quality of human presence into a measurable economic variable. By benchmarking output quality against AI, it ensures that the human premium is earned rather than assumed. By settling compensation on-chain with immutable inputs, it gives workers a tamper-proof record of their contribution that no employer can retroactively alter. And by setting a guaranteed floor exchange rate, it ensures that verified human work always has a minimum dollar value, regardless of market fluctuations. The numbers in this paper -- $3.00 per hour for unskilled work, $118.80 for expert professional work, all grounded in physiological verification -- are not projections. They are what the formula produces given calibrated default parameters. They are adjustable by governance. But they are honest: they reflect what it costs to have a verified human being, at each level of expertise, genuinely present and engaged. That is what we are paying for. The HPT framework makes sure we actually pay for it.

REFERENCES

Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives, 33*(2), 3–30.

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from U.S. labor markets. *Journal of Political Economy, 128*(6), 2188–2244.

Benedek, M., & Kaernbach, C. (2010). A continuous measure of phasic electrodermal activity. *Journal of Neuroscience Methods, 190*(1), 80–91.

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings, 108*, 43–47.

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). *Sparks of artificial general intelligence: Early experiments with GPT-4* (arXiv preprint arXiv:2303.12528).

Buterin, V. (2014). *A next-generation smart contract and decentralized application platform* (White paper).

Castaldo, R., Melillo, P., Bracale, U., Caserta, M., Triassi, M., & Pecchia, L. (2021). Heart rate variability analysis as a diagnostic marker for autonomic dysfunction. *Frontiers in Physiology, 12*, Article 724218.

Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies, 2*(3), 200–219.

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794).

Fairclough, S. (2009). Fundamentals of physiological computing. *Interacting with Computers, 21*(1–2), 133–145.

Halbe, S. (2024). *Presence cannot be faked: Detecting conscious human engagement versus AI impersonation through cross-modal biosignal coherence* (Manuscript under review).

Healey, J. A., & Picard, R. W. (2005). Detecting stress during real-world driving tasks using physiological sensors. *IEEE Transactions on Intelligent Transportation Systems, 6*(2), 156–166.

International Labour Organization. (2021). *World employment and social outlook: The role of digital labour platforms in transforming the world of work*.

Kreibig, S. D. (2010). Autonomic nervous system activity in emotion: A review. *Biological Psychology, 84*(3), 394–421.

Mucha, A., Khalid, M., & Tetzlaff, R. (2020). Cross-modal coherence in physiological signals as a marker of attentional engagement. *Frontiers in Human Neuroscience, 14*, Article 116.

Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*.

Thayer, J. F., & Lane, R. D. (2000). A model of neurovisceral integration in emotion regulation and dysregulation. *Journal of Affective Disorders, 61*(3), 201–216.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In I. Guyon et al. (Eds.), *Advances in Neural Information Processing Systems* (Vol. 30, pp. 5998–6008).

Webb, M. (2020). *The impact of artificial intelligence on the labor market* (SSRN Working Paper No. 3482150).

World Economic Forum. (2023). *The future of jobs report 2023*.

Conflict of Interest
The author(s) declare that there are no conflicts of interest regarding the publication of this manuscript.

Funding
This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contributions
The author(s) contributed substantially to the conception and design of the study, data collection, data analysis, and manuscript preparation. All authors reviewed and approved the final version of the manuscript.

Informed Consent
Informed consent was obtained from all participants involved in the study.

Use of Generative AI
The author(s) confirm that generative AI tools were used solely for minor language refinement purposes and did not contribute to the intellectual content, analysis, interpretation, or conclusions of the study.